EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	1745	cost$1 optimiz$5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:48
L2	2334	optimiz$5 cost$1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:48
L3	61	(L1 or L2) and join and table$2 and unique and sum	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:48
L4	451	merge join	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:49
L5	0	L1 and L2 and L4 and(confidence level)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:49
L6	0	(optimiz$5 cost$1) and (unoptimiz$5 cost$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:50
L7	0	(optimiz$5 cost$1) and (unoptimiz$5 cost$1) and join	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:50
L8	60	(optimiz$5) and (unoptimiz$5) and join and cost and table	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:52

EAST Search History

L9	8	(optimiz$5) and (unoptimiz$5) and join and cost and table and (optimiz$5 near3 cost)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:51
L10	0	(optimiz$5) and (unoptimiz$5) and join and cost and table and (optimiz$5 near3 cost) and (unoptimiz$5 near3 cost)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 13:51
L13	0	((optimiz$5) and (unoptimiz$5) and join and cost and table and quer$4).clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	ADJ	ON	2007/05/08 14:53

Google Scholar BETA

Web Images Video News Maps **more »**

join and "optimized cost" and "unoptimized co Search

Advanced Scholar Search
Scholar Preferences
Scholar Help

The **"AND"** operator is unnecessary -- we include all search terms by default. [details]

Tip: Try removing quotes from your search to get more results.

Your search - **join and "optimized cost" and "unoptimized cost"** - did not match any articles.

Suggestions:

- Make sure all words are spelled correctly.
- Try different keywords.
- Try more general keywords.
- Try fewer keywords.
- Try your query on the entire web.

Google Home - About Google - About Google Scholar

©2007 Google

Sign in

Google

Web Images Video News Maps **more »**

join and "optimized cost" and "unoptimized co Search Patents

Advanced Patent Search
Google Patent Search Help

The **"AND"** operator is unnecessary -- we include all search terms by default. [details]

Patents

(**0.01** seconds)

« View all web results for **join and "optimized cost" and "unoptimized cost"**

Your search - **join and "optimized cost" and "unoptimized cost"** - did not match any documents.

Suggestions:

- Make sure all words are spelled correctly.
- Try different keywords.
- Try more general keywords.
- Try fewer keywords.
- Try your search at www.uspto.gov

join and "optimized cost" and "unoptimized co Search Patents

Google Patent Search Help | Advanced Patent Search

Google Home - About Google - About Google Patent Search

©2007 Google

Sign in

Google

Web Images Video News Maps **more »**

join and "optimized cost" and "unoptimized co Search Advanced Search Preferences

The **"AND"** operator is unnecessary -- we include all search terms by default. [details]

Web

Tip: Try removing quotes from your search to get more results.

Your search - **join and "optimized cost" and "unoptimized cost"** - did not match any documents.

Suggestions:

- Make sure all words are spelled correctly.
- Try different keywords.
- Try more general keywords.
- Try fewer keywords.

Google Home - Advertising Programs - Business Solutions - About Google

©2007 Google

PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: ● The ACM Digital Library ○ The Guide

+cost and +join and query and +("optimized cost") and +("uno

SEARCH

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used **cost** and **join** and **query** and **optimized cost** and **unoptimized cost** and **table** and **unique** and **compare** and **duplicate**

Found **29** of **201,062**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 29

Result page: **1** 2 next

Relevance scale

1 Accelerating XPath evaluation in any RDBMS

Torsten Grust, Maurice Van Keulen, Jens Teubner

March 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 1

Publisher: ACM Press

Full text available: pdf(781.01 KB) Additional Information: full citation, appendices and supplements, abstract, references, cited by, index terms

This article is a proposal for a database index structure, the *XPath accelerator*, that has been specifically designed to support the evaluation of XPath path expressions. As such, the index is capable to support *all* XPath axes (including ancestor, following, preceding-sibling, descendant-or-self, etc.). This feature lets the index stand out among related work on XML indexing structures which had a focus on the child and descendant axes only. The index has been designed with a close ...

Keywords: Main-memory databases, XML, XML indexing, XPath

2 Research sessions: Research 26: Query optimization: Analysis of two existing and one new dynamic programming algorithm for the generation of optimal bushy join trees without cross products

Guido Moerkotte, Thomas Neumann

September 2006 **Proceedings of the 32nd international conference on Very large data bases - Volume 32 VLDB'2006**

Publisher: VLDB Endowment

Full text available: pdf(599.73 KB) Additional Information: full citation, abstract, references, index terms

Two approaches to derive dynamic programming algorithms for constructing join trees are described in the literature. We show analytically and experimentally that these two variants exhibit vastly diverging runtime behaviors for different query graphs. More specifically, each variant is superior to the other for one kind of query graph (chain or clique), but fails for the other. Moreover, neither of them handles star queries well. This motivates us to derive an algorithm that is superior to the t ...

3 Tracing the lineage of view data in a warehousing environment

Yingwei Cui, Jennifer Widom, Janet L. Wiener

June 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(458.90 KB) Additional Information: full citation, abstract, references, citings, index terms

We consider the view data lineageproblem in a warehousing environment: For a given data item in a materialized warehouse view, we want to identify the set of source data items that produced the view item. We formally define the lineage problem, develop lineage tracing algorithms for relational views with aggregation, and propose mechanisms for performing consistent lineage tracing in a multisource data warehousing environment. Our result can form the basis of a tool that al ...

Keywords: data warehouse, derviation, lineage, materialized views

4 Dynamic data fusion for future sensor networks

Umakishore Ramachandran, Rajnish Kumar, Matthew Wolenetz, Brian Cooper, Bikash Agarwalla, Junsuk Shin, Phillip Hutto, Arnab Paul
August 2006 **ACM Transactions on Sensor Networks (TOSN)**, Volume 2 Issue 3
Publisher: ACM Press
Full text available: pdf(2.44 MB) Additional Information: full citation, abstract, references, index terms

DFuse is an architectural framework for dynamic application-specified data fusion in sensor networks. It bridges an important abstraction gap for developing advanced fusion applications that takes into account the dynamic nature of applications and sensor networks. Elements of the DFuse architecture include a fusion API, a distributed role assignment algorithm that dynamically adapts the placement of the application task graph on the network, and an abstraction migration facility that aids such ...

Keywords: Sensor network, data fusion, energy awareness, in-network aggregation, middleware, platform, role assignment

5 Garbage collecting the Internet: a survey of distributed garbage collection

Saleh E. Abdullahi, Graem A. Ringwood
September 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 3
Publisher: ACM Press
Full text available: pdf(337.65 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Internet programming languages such as Java present new challenges to garbage-collection design. The spectrum of garbage-collection schema for linked structures distributed over a network are reviewed here. Distributed garbage collectors are classified first because they evolved from single-address-space collectors. This taxonomy is used as a framework to explore distribution issues: locality of action, communication overhead and indeterministic communication latency.

Keywords: automatic storage reclamation, distributed, distributed file systems, distributed memories, distributed object-oriented management, memory management, network communication, object-oriented databases, reference counting

6 Migration: The design and implementation of Zap: a system for migrating computing environments

Steven Osman, Dinesh Subhraveti, Gong Su, Jason Nieh
December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI
Publisher: ACM Press